

This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# Knowledge Bridger: Towards Training-Free Missing Modality Completion

Guanzhou Ke<sup>1,2</sup>, Shengfeng He<sup>2,\*</sup>, Xiaoli Wang<sup>6</sup>, Bo Wang<sup>4,\*</sup>,

Guoqing Chao<sup>5</sup>, Yuanyang Zhang<sup>3</sup>, Yi Xie<sup>7</sup>, Hexing Su<sup>8</sup>

<sup>1</sup>Beijing Jiaotong University, <sup>2</sup>Singapore Management University, <sup>3</sup>Southeast University,

<sup>4</sup>Institute of Automation, Chinese Academy of Sciences, <sup>5</sup>Harbin Institute of Technology,

<sup>6</sup>Nanjing University of Science and Technology, <sup>7</sup>South China University of Technology, <sup>8</sup>Xiamen Institute of Technology

### Abstract

Previous successful approaches to missing modality completion rely on carefully designed fusion techniques and extensive pre-training on complete data, which can limit their generalizability in out-of-domain (OOD) scenarios. In this study, we pose a new challenge: can we develop a missing modality completion model that is both resourceefficient and robust to OOD generalization? To address this, we present a training-free framework for missing modality completion that leverages large multimodal model (LMM). Our approach, termed the "Knowledge Bridger", is modality-agnostic and integrates generation and ranking of missing modalities. By defining domain-specific priors, our method automatically extracts structured information from available modalities to construct knowledge graphs. These extracted graphs connect the missing modality generation and ranking modules through the LMM, resulting in high-quality imputations of missing modalities. Experimental results across both general and medical domains show that our approach consistently outperforms competing methods, including in OOD generalization. Additionally, our knowledge-driven generation and ranking techniques demonstrate superiority over variants that directly employ LMMs for generation and ranking, offering insights that may be valuable for applications in other domains.

## 1. Introduction

Multimodal learning [17, 33] is a foundational task in artificial intelligence that enables models to integrate diverse data types—such as text, images, and audio—to enhance their representational and reasoning capabilities [27, 28, 40]. However, real-world applications often encounter incomplete or missing modalities due to factors like noise, sensor failures, or privacy constraints, which

can compromise a model's robustness and generalizability. Missing modality completion (MMC) addresses this issue by imputing or reconstructing absent modalities based on available data, allowing models to fully leverage multimodal information. Recent MMC methods [3, 13, 21, 49] have demonstrated promising results across various applications. For instance, in medical diagnostics [6, 55], MMC models can reconstruct absent diagnostic data (e.g., MRI, CT, and X-rays) by leveraging existing medical reports or related images.

Despite these advancements, MMC methods often have limited out-of-domain (OOD) transferability, typically requiring retraining on new domain data to maintain performance. For example, some MMC models [30, 57] developed for sentiment analysis need substantial adaptations to be effective in applications like medical imaging or autonomous driving, leading to significant human and computational costs. To address these challenges, recent works have proposed domain-agnostic solutions, such as missing modality tags [54] to aid in predicting absent modalities or prompt-learning techniques [13, 21] that dynamically adjust fusion strategies. While these methods lower domain adaptation costs, they still require extensive training data, limiting their effectiveness in data-scarce domains, such as rare disease analysis. This raises a crucial question: Can we develop an MMC model that achieves both low-resource dependency (e.g., minimal computational and domain-specific data requirements) and strong OOD capability?

Recent advancements [16, 26, 44, 48] in large multimodal models have highlighted their strong OOD capabilities and adaptability to new tasks with minimal resources [20]. In this work, we explore the potential of leveraging LMM to effectively address the MMC challenge. Before progressing, it is essential to distinguish between modality generation and missing modality completion. Modality generation typically focuses on creating new modalities from available information, prioritizing **diversity and creativity** in the generated content. In contrast, missing modality completion emphasizes **accuracy** over diversity, recon-

<sup>\*</sup>Corresponding authors (Email: shengfenghe@smu.edu.sg, wangbo@ia.ac.cn).

structing absent modalities to maintain semantic coherence and improve task performance.

Addressing MMC with LMM presents two primary challenges: 1) generation: applying constraints on modality generation to ensure fidelity to the missing content, and 2) ranking: selecting the most appropriate completion from generated candidates. Effective use of LMM for these tasks requires embedding sufficient domain knowledge to guide accurate interpretation and reconstruction of missing content. For example, generating an accurate image in a vision-language task can be challenging with only brief textual descriptions; however, incorporating detailed descriptions of entities and their interactions allows LMM to more accurately reconstruct missing data. Additionally, such prior knowledge aids LMM in ranking candidate completions by identifying the most semantically plausible options. Leveraging LMM's in-context learning capability [52], we can non-intrusively embed prior knowledge, such as using the Chain-of-Thought (CoT) approach [53], to enhance the model's understanding of missing modalities.

Based on these insights, we propose a novel, trainingfree MMC approach, termed "Knowledge Bridger", which autonomously mines multimodal knowledge, generates missing modalities, and ranks the best completions. This method comprises three main modules: a knowledge modeling module, a knowledge-driven modality generation module, and a ranking module. In the knowledge modeling stage, we employ LMM to analyze available modalities and extract key elements such as objects, interactions, and attributes using the CoT approach. For specialized fields like medical imaging, we use in-context learning to integrate domain-specific knowledge, reducing reliance on extensive target domain data. Following knowledge extraction, we construct a knowledge graph to represent the relationships and attributes of available and missing modalities, providing LMM with a structured reference for generating missing content. To ensure high-quality completion, the knowledgedriven generation module utilizes this graph to guide LMM in generating specific content with precise details, such as object locations and observable attributes. Finally, the ranking module, informed by expert knowledge, assesses each generated completion by computing graph and representational similarity scores. Specifically, a graph similarity score is derived from comparing the knowledge graphs of the available and generated modalities, while a representational similarity score is calculated using models like BLIP [22] and CLIP [38]. A weighted average of these scores offers the final assessment.

Our extensive experiments indicate that our method markedly improves MMC performance in both general and OOD scenarios. We also find that our approach scales effectively with LMM, with larger models yielding higherquality completions. For instance, using OpenAI's GPT-40 [16] results in a significant performance boost across metrics compared to models with 72B or 7B parameters. The visualization results show that our method outperforms the conditional generation variant in generating missing modalities. Our contributions are summarized as follows:

- We introduce a training-free pipeline to address missing modality completion, leveraging LMM to automatically extract multimodal knowledge, generate missing modalities, and rank completions. To our knowledge, this is the first work to apply LMM to MMC tasks.
- We delve into a modality-agnostic unified strategy for both missing modality completion and ranking. This approach allows us to focus on defining domain-specific knowledge without the necessity of intricate fusion methods or implementing a specialized training pipeline.
- We present extensive experimental evidence demonstrating that our method facilitates domain transfer, outperforming other MMC methods in both general and OOD scenarios. Additionally, our generated completion data improves the performance of other MMC models.

### 2. Method

### 2.1. Overview

Our approach aims to construct a training-free MMC pipeline by leveraging a pre-trained LMM. This pipeline extracts and models knowledge from available modalities, and subsequently uses this knowledge to generate missing modality data and select the most appropriate generated candidate. The pipeline is depicted in Fig. 1 and consists of three steps:

- Step 1: construction of a knowledge graph from available modalities (in section 2.2). The goal of this step is to utilize the general knowledge a pre-trained LMM to understand the content of each modality and their interrelationships.
- Step 2: knowledge-driven generation (in section 2.3). In this step, the LMM employs the knowledge graph to extract specific details about the missing modality, including the number and attributes of objects. A corresponding modality generator then uses this information to generate the required missing modality.
- Step 3: knowledge-based ranking (in section 2.4). This step aims to compute a quality score for the generated missing data by assessing the graph similarity and semantic similarity between the missing and available modalities.

#### 2.2. Knowledge Graph Modeling

Our objective is to automatically extract knowledge from available modalities to support the generation and ranking of any missing modalities. In our context, 'knowledge' refers to information encapsulating the characteristics



Figure 1. **Overview of 'Knowledge Bridger' Pipeline.** The pipeline consists of three steps: 1) construction of a knowledge graph from available modalities, 2) knowledge-driven generation of missing modalities, and 3) knowledge-based ranking. The input of this pipeline is the available multimodal data and known domain knowledge. The output is the required missing modality.

of existing modalities, enabling the generation and ranking modules to create semantically consistent missing data. Extracting relevant knowledge from unknown domains, however, presents significant challenges. From a knowledge graph perspective, constructing a meaningful, modality-specific graph requires predefined nodes and relationships. In a training-free context, predefining these elements becomes particularly challenging.

To overcome this, we develop an automatic entity and relationship mining strategy using an LMM. This strategy leverages the extensive prior knowledge and OOD capabilities of LMM to identify entities and relationships across various modalities, even without predefined elements. Recent study [20] highlight LMM's potential for zero-shot learning and reasoning.

Building on this analysis, LMM can be utilized to extract elements from available modalities using prompts. To enhance scalability, we propose the following extraction rule: *[Entity: Reasoning Prompts]*. For instance, to identify potential objects, we could use: *('Objects': 'Identify the major objects in the [modality-type].')*. This rule allows us to incorporate object relationships and interaction data<sup>1</sup>. To improve adaptability across domains, domainspecific prior knowledge—such as histological and clinical diagnostic information for medical image analysis—can be included. This approach offers two main advantages: 1) reducing misconceptions when LMM operates in a new domain, and 2) enhancing its reasoning capabilities for novel

#### entities.

LMM can further consolidate the extracted information into a modality-specific knowledge graph. A straightforward approach is to guide LMM to extract potential entityrelation pairs from the collected data. However, this method is constrained by the context window length, and an excessive number of extraction rules may lead to overlooked entity-relation pairs. To mitigate these limitations, we implement the Chain-of-Thought (CoT) method. Specifically, LMM is first directed to produce concise responses for each rule, followed by extracting unique entity-relation pairs from these responses. This stepwise decomposition improves both the accuracy of responses for each rule and the synthesis of information. Importantly, we extract and retain data solely from current modalities to avoid interference from unrelated information, thereby enhancing LMM's reasoning efficiency.

#### 2.3. Knowledge-driven Generation

The goal of missing multimodal generation is to understand the content in available modalities and generate missing ones that align semantically. Two critical factors impact the quality of the generated missing modalities: understanding multimodal content and maintaining consistency. We previously discussed using LMM for content comprehension and knowledge graph extraction. Here, we explore using LMM for ensuring consistency and guiding generation. For convenience, we use image-text pairs as our study objects. For example, when an image is available, we aim

<sup>&</sup>lt;sup>1</sup>The complete rules and prompts are provided in the appendix.

to generate text that closely matches real data. A basic approach is using LMMs to describe the image directly. However, this method involves significant randomness. First, the form of the missing text is unknown—it could be a caption, summary, or description. Second, we cannot precisely specify the subject of the missing text.

To solve this, we propose a knowledge-driven entity alternation strategy. Using domain knowledge and extracted knowledge graphs, we select related entities.For instance, if missing data focuses on an entity such as an 'objection,' we traverse elements within the knowledge graph related to 'objection.' We then adopt a multi-view generation manner, allowing the LMM to generate missing information with each element as the subject, while encompassing all nodes and attributes within the knowledge graph. These outputs are stored as standardized text descriptions, reducing randomness, enhancing result retrieval, and offering better control and interpretability. With these descriptions, modality generators can create missing data. For missing images, entity-based descriptors can guide conditional diffusion methods. For missing text, LMM process these descriptions to generate outputs. This approach works across various fields, using mature generation models with domain knowledge to create needed data without extra training. However, relying only on this may not guarantee full accuracy, which we discuss in the next section.

#### 2.4. Knowledge-based Ranking

To achieve automatic ranking of generated missing data based on the provided knowledge, we introduce graph similarity and representation similarity. Graph similarity is computed as the average cosine similarity score of the adjacency matrices of two graphs, as presented in the following equation:

$$\cos_{graph}(A_i, B_i) = \frac{1}{n} \sum_{i=1}^{n} \frac{A_i \cdot B_i}{\|A_i\| \|B_i\|},$$
 (1)

where  $A_i$  and  $B_i$  represent the *i*-th row vectors of adjacency matrices A and B, respectively.  $||A_i||$  refers to the Euclidean norm of the *i*-th row:  $||A_i|| = \sqrt{\sum_{j=1}^m A_{ij}^2}$ . n and m represent the rows and columns of the adjacency matrix, respectively. This metric reflects the degree of similarity between the two graphs, and its value is normalized between 0 and 100. On the other hand, we compute the representation similarity between the generated and available modalities to reflect semantic consistency. For consistency, we similarly employ cosine similarity to compute the similarity between two representations, expressed as  $\cos(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}$ , where  $\mathbf{a}$  and  $\mathbf{b}$  are the vectors of two modalities. Inspired by [2], we directly utilize CLIP [38] and BLIP [22] to obtain semantic embeddings for each modality. Finally, we derive the following equation to com-

pute the generation quality score between any pair of available and missing modalities:

$$QS(x_a, x_m) = \cos_{graph}(f_a(x_a), f_a(x_m)) + [\cos(f_c(x_a), f_c(x_m)) + \cos(f_b(x_a), f_b(x_m))],$$
(2)

where  $x_a$  and  $x_m$  represent available and missing modalities, respectively. The functions  $f_a(\cdot)$ ,  $f_c(\cdot)$ , and  $f_b(\cdot)$  are used to obtain the adjacency matrix, CLIP's embedding, and BLIP's embedding of the given modality, respectively. We argue that  $QS(\cdot, \cdot)$  can comprehensively assess two critical factors: knowledge structure similarity and semantic consistency. A higher value indicates a higher quality of the generated missing modality. Our method ultimately outputs the generated missing modality with the highest score.

### **3. Experiments**

### 3.1. Setup

Baselines. To validate the effectiveness of our proposed method, we selected several baseline methods, categorized into two types: imputation-based and non-imputationbased. Imputation-based methods aim to restore missing modalities by learning relationships among the representations within each modality; for this category, we choose MMIN [57] and DiCMoR [49] as baselines. In contrast, non-imputation-based methods bypass the need to restore missing modalities or predict their representations, allowing downstream tasks to be completed without these steps. We choose MPLMM [13] and MPMM [21] as baselines for this category. Additionally, we introduce a simple baseline that completely removing missing data. This baseline utilizes the pre-trained CLIP model as multimodal backbones. Subsequently, the modality features are concatenated and fed into a classification head (a single-layer MLP) to ultimately obtain the probability for each category.

Dataset. To evaluate the domain transferability of our method, we employ two general domain datasets and one OOD dataset. The general domain datasets include COCO-2014 [25] and MM-IMDb [1]. COCO-2014 [25] is a large-scale vision-language dataset containing approximately 81K training samples and 41K validation and test samples, with objects classified into 80 categories. In our setup, we treat it as a multimodal multi-label classification dataset, using the validation set as the test set. MM-IMDb [1] is a large-scale movie genre classification dataset with around 25,000 movies from IMDb, each represented by its movie poster and plot summary, and annotated with 27 multi-label genre tags. In our setup, we divide the dataset into 18,160 samples for the training set and 7,799 samples for the test set. On the other hand, we employ the IU X-ray [7] dataset as our OOD dataset. This dataset includes 6,674 training and 756 testing samples. We further expanded the original 14 observations into 105 finer-grained categories based on location and severity, creating a long-tail dataset. **Evaluation Metrics.** For multimodal multi-label classification on three datasets, we report the performance of comparison methods using the macro F1-score and Average Precision (AP). We report the average result of five different random seeds. Additionally, we introduce the mean similarity score (denoted **SS**) to evaluate the generation quality of our method and other imputation-based methods. Specifically, we utilize the vision and text backbones of the pretrained CLIP [38] to compute embeddings for all groundtruth and generated missing modalities. The cosine similarity between each pair of embeddings is then computed and averaged, with scores ranging from 0 to 100. A higher score indicates better generation quality.

Implementation Details. We employ Qwen2-VL-7B [48] as our default large multimodal model<sup>2</sup>. It is specifically designed to handle both visual and textual inputs, enabling it to process and interpret complex multimodal content. For image reconstruction, we apply Stable Diffusion XL (SDXL) 1.0 [36] as the restoration module for general domains. SDXL 1.0 is an advanced text-to-image diffusion model that can generate images according to a given prompt. Additionally, for the restoration of chest X-ray modality, we use Cheff [51], a cascaded chest X-ray latent diffusion model. By default, we generate 5 candidates for the missing modality during the generation process. Since the chosen LMM has already been pre-trained on the COCO dataset, to ensure a fair comparison, we employ pre-trained CLIP [38] vision and text encoders to replace the modality backbones of the comparison methods. We conduct all experiments on the PyTorch 2.4.0 platform, running on Ubuntu 20.04 LTS utilizing 4 GPUs (NVIDIA GeForce RTX 4090 with 24 GB of memory). In our setting, we conduct the missing rate  $\eta = \{0.3, 0.5, 0.7\}$  to simulate the missing modality scenario during training.

#### 3.2. Quantitative Analysis

In Tables 1 and 2, we present the results of our method compared to others across different missing data ratios, as measured by F1, AP, and mean similarity score (SS) on three datasets <sup>3</sup>. To comprehensively evaluate the performance of various approaches, we introduce two baselines: one trained on complete data and another trained on data with missing entries removed. Additionally, to fair comparison regarding the scale of model parameters, the proposed method based on the Qwen-VL-2B were incorporated.

Results from general domain datasets (COCO-2014 and MM-IMDb) indicate that most MMC methods outperform the baseline where missing data is simply removed. Our method demonstrates superior performance across different

missing rates, particularly at a missing rate of 0.7, where it shows significant improvements in both F1 and AP metrics. We attribute these enhancements to the synthetic data generated by our method, which appears to bolster downstream task performance. This phenomenon is also observed in [12, 47]. However, it should be noted that the synthetic data may lead to a slight decline in mean similarity score.

On the other hand, results from the OOD dataset highlight our method's superior domain adaptation capability compared to other methods. Our approach not only excels in classification metrics but also significantly surpasses other imputation-based methods, such as MMIN and DiC-MoR, in missing data generation metrics (SS). This advantage can be credited to our method's ability to harness the few-shot learning and in-context learning capabilities of LMM, effectively mitigating the impact of varying degrees of data incompleteness.

### 3.3. Ablation Study

Q1: Does 'knowledge' really help MMC? We study the impact of different components under general and OOD scenarios with the same missing rate. As shown in Table 3, we report the results of various components. For the generation process, we analyze the variant without knowledge modeling<sup>4</sup> (row 1) and the variant with random ranking (row 2). Additionally, in examining the ranking module, we assess the variant with random ranking (row 3), the variant without knowledge graph ranking (row 4), and the variant without using semantic similarity score ranking (row 5). The results demonstrate that the variant without knowledge modeling shows a significant decline across all metrics under OOD scenarios. Furthermore, the semantic similarity ranking strategy in the ranking module is crucial for the effectiveness of the MMC. In general, the study reveals that knowledge modeling is the most critical component of our method.

**Q2:** Does the scale of model parameters affect MMC? We evaluate the performance of our approach using different model scales under at various missing rates. Specifically, we conduct the 2B, 7B, and 72B variants of the opensource large model Qwen-VL [48], as well as OpenAI's GPT-40 [16] on the IU X-ray dataset. For a fair comparison, we leverage these models solely for handling the knowledge modeling and integration in steps 1 and 2 of our method, keeping the modality generators unchanged. As illustrated in Fig. 2, GPT-40 demonstrates significant superiority and robustness. These results demonstrate that as the scale of model parameters increases, the quality of knowledge modeling in our approach also improves.

Q3: Does generating more missing data improve performance? We evaluate the impact of different numbers of

<sup>&</sup>lt;sup>2</sup>More implementation details are presented in the appendix.

<sup>&</sup>lt;sup>3</sup>More results (more modalities) are presented in the appendix.

<sup>&</sup>lt;sup>4</sup>We directly use LMM to generate descriptions of available modalities.

		COCO-2014 [25]						MM-IMDb [7]										
Missing Rate $\eta$		0.3			0.5			0.7			0.3			0.5			0.7	
Method	F1	AP	SS	F1	AP	SS	F1	AP	SS	F1	AP	SS	F1	AP	SS	F1	AP	SS
Baseline (complete) Baseline (remove missing)	75.8	F 80.8	1: 78.   -	3   AI  73.1	P: 84.0	5   SS   -	: -  70.3	77.5	-	51.8	F  58.5	1: 56. -	2   AI   50.3	2: 62.2 56.1	7   SS   -	: -  47.2	53.9	-
MPMM [21] (CVPR'23) MPLMM [13] (ACL'24) MMIN [57] (ACL'21) DiCMoR [49] (CVPR'23)	76.2 <u>77.1</u> 73.2 65.3	82.0 82.6 78.3 74.4	$\frac{-}{37.1}$ $\frac{37.1}{34.3}$	74.7 75.2 71.4 59.7	80.5 81.3 77.5 67.1	$\frac{-}{36.7}$	71.0 72.3 70.5 55.3	78.8 80.1 76.4 64.0	$\begin{vmatrix} -\\ -\\ 36.0\\ 31.9 \end{vmatrix}$	53.6 53.9 50.1 49.2	59.8 60.3 53.8 54.7	- 26.7 <u>26.9</u>	51.1 52.8 49.5 43.7	56.9 57.3 51.7 50.8	$\frac{\frac{-}{26.3}}{25.9}$	$ \begin{array}{r}     48.5 \\     \underline{49.1} \\     44.6 \\     30.5 \\   \end{array} $	55.7 <u>56.2</u> 50.8 41.7	$\frac{\frac{-}{2}}{\frac{24.4}{23.1}}$
Ours (Qwen-VL-2B) Ours (Qwen-VL-7B)	76.1 <b>77.5</b>	81.8 <b>82.9</b>	39.8 <b>40.4</b>	74.3 77.5	79.9 <b>82.8</b>	37.1 <b>40.2</b>	70.7 77.9	78.2 <b>83.5</b>	36.3 38.2	52.4 <b>54.7</b>	58.9 <b>60.9</b>	28.8 <b>33.5</b>	51.9 <b>54.9</b>	57.1 <b>61.3</b>	28.3 <b>32.7</b>	50.3 55.2	56.5 <b>61.8</b>	28.1 <b>32.3</b>
$\Delta$ Complete Baseline $\Delta$ SOTA	-2.8 +0.4	-1.7 +0.3	- +3.3	-0.8 +2.3	-1.8 +1.5	- +3.7	-0.4 +5.6	-1.1 +3.4	+2.5	-1.5 +0.8	-1.8 +0.6	- +6.6	-1.3 +2.1	-1.4 +4.0	- +6.4	-1.0 +6.1	-0.9 +5.6	- +7.9

Table 1. Quantitative analysis results (%) on COCO-2014 and MM-IMDb datasets. Bold denotes the best results and <u>underline</u> denotes the second-best. SS (%) refers to the average similarity score, which is used to assess the generation quality of imputation-based methods. A higher score indicates better quality. '-' indicates that the metric is not applicable. All results are reproduced using the officially released code.

Missing Rate $\eta$	0.3	0.7			
Method	F1   AP   SS	F1   AP   SS			
Baseline (complete) Baseline (remove missing)	F1: 57.0   A 49.1   71.4   -	P: 75.7  SS: -   31.5   56.2   -			
MPMM [21] (CVPR'23) MPLMM [13] (ACL'24) MMIN [57] (ACL'21) DiCMoR [49] (CVPR'23)	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$ \begin{vmatrix} \underline{36.8} \\ 35.2 \\ 26.7 \\ 29.8 \\ 53.6 \end{vmatrix} \begin{vmatrix} 61.4 \\ - \\ - \\ 29.8 \\ 53.6 \\ \underline{13.3} \end{vmatrix} $			
Ours (Qwen-VL-2B) Ours (Qwen-VL-7B)	51.4 72.0 21.9 53.6 73.9 22.6	41.1     68.9     17.7       46.3     70.5     19.8			
$\Delta$ Complete Baseline $\Delta$ SOTA	-3.4 -1.8 - +3.7 +1.2 +4.5	-10.7 -5.2 - +9.5 +8.6 +6.5			

Table 2. Quantitative results (%) on IU X-ray datasets. Bold denotes the best results and <u>underline</u> denotes the second-best. SS (%) refers to the average similarity score, which is used to assess the generation quality of imputation-based methods. A higher score indicates better quality. '-' indicates that the metric is not applicable. All results are reproduced using the officially released code.

generation candidates in step 2 of our method on the MM-IMDb and IU X-ray datasets. Specifically, we conduct the settings with generation candidates  $n = \{1, 5, 10, 15\}$ , and the results are presented in Table 4. The findings indicate that generating only one candidate significantly reduces the missing generation quality. While using a higher number of generation candidates can lead to slight improvements across all metrics, it also substantially increases inference time. Therefore, setting n = 5 provides a balance between inference time and performance.

Missing Rate $\eta$	MM-IMDb 0.7	IU X-ray 0.7			
Variants	F1   AP   SS	F1   AP   SS			
0 Baseline (Qwen-VL-7B)	55.2 61.8 32.3	46.3   70.5   19.8			
1 <i>w/o</i> Knowledge Modeling 2 + Random Ranking	-1.3  -3.6  -8.8 -1.6  -4.1  -9.9	-17.5  -29.2  -13.7 -19.3  -31.8  -15.0			
<ul> <li>3 Random Ranking</li> <li>4 w/o Knowledge Ranking</li> <li>5 w/o Semantic Ranking</li> </ul>	-0.5-2.7-0.6-0.2-0.8-0.1-0.2-1.0-0.3	-3.8-7.1-4.7-1.9-2.7-1.1-2.4-3.3-1.6			

Table 3. **The impact of various components.** We report the comparison results between different combinations and the baseline.

Missing Rate $\eta$	M	M-IM 0.7	Db	IU X-ray 0.7			
Candidates $n$	time	F1	AP	SS	F1	AP	SS
n = 5 (default)	40s	55.2	61.8	32.3	46.3	70.5	19.8
n = 1	7s	-0.7	-2.8	-7.3	-8.5	-16.4	-14.7
n = 10	83s	+0.4	+0.9	+0.1	+0.0	+0.1	+0.0
n = 15	132s	+0.4	+0.9	+0.1	+0.1	+1.3	+0.3

Table 4. The impact of different generation candidates.

#### 3.4. Visualization Analysis

To better understanding the differences between our approach and conditional generation, we present completion results produced by our method, as shown in Fig. 3. We provide visualization results across different missing modalities for three datasets<sup>5</sup>. In the general domain, our knowledge modeling module focuses on understanding the quantity of objects, their attributes, and the contextual en-

<sup>&</sup>lt;sup>5</sup>More visualization are presented in appendix.



(b) Average similarity score at different missing rates.

Figure 2. The impact of different model's parameters scales and missing rates.

vironment. The results in Fig. 3a and Fig. 3b indicate that our method is more similar with the original missing modality compared to direct generation approaches. In the medical domain, incorporating knowledge of different lesions enables the LMM to understand the relationships between various regions in chest X-rays and the content described by the modality. Fig. 3c and Fig. 3d show that our method provides a more reliable strategy for missing data completion than direct generation.

### 4. Related Work

## 4.1. Missing Multimodal Learning

Missing multimodal learning enables models to train and inference effectively even when some modalities are absent. The missing multimodal learning methods can be categorized into imputation-based and non-imputation-based approaches. Imputation-based methods predict or recover missing data by learning interactions between modalities. Typical strategies include filling with special values such as zero or the average value, which are unsuitable for highdimensional data [34]. Advanced methods like SMIL [30] use meta-learning to simulate missing modalities. Wang *et al.* [46] propose a multitask framework for missing modality learning, which learns shared features across different tasks and uses these shared features to reconstruct missing modalities. Lian *et al.* [24] use graph neural networks to explore relationships between modalities and recover missing ones, while Zhao *et al.* [57] and Pham *et al.* [35] leverage cycle consistency to capture cross-modal interactions for predicting missing modalities.

On the other hand, non-imputation-based methods focus on fusion strategies or missing indicators instead of predicting missing data [18]. Wang *et al.* [50] and Han *et al.* [15] use translational semantics for implicit fusion. Ma *et al.* [31] leverage transformers' variable input capability. Lee *et al.* [21] and Guo *et al.* [13] introduce learnable missing semantic tokens, while Zeng *et al.* [54] use fixed tags to identify missing inputs. Current methods often require pre-training with complete data and show limited domain transfer. Our approach leverages LMM to predict missing modalities using domain knowledge without extra pre-training, employing few-shot learning for flexible adaptation.

#### 4.2. Large Multimodal Models

Large multimodal models (LMMs) integrate diverse data types, such as text, images, and audio, to enhance understanding and interaction capabilities. Early works focused on fusion techniques to combine these modalities, like early [4], late [9], or hybrid fusion [37]. Recent advancements [44, 45], such as OpenAI's GPT-40 [16] and Google's PaLM-E [8], have achieved remarkable success by increasing the scale of both pre-training data and model parameters. These private ultra-models (greater than 100 billion parameters) exhibit robust generalization across various domains, which enables applications in image captioning, visual question answering, and multimodal dialogue. On the other hand, some open-source, medium- to small-scale (less than 100 billion parameters) models, such as LLaVA [26], Qwen-VL [48], and InternVL [5], have matched the performance of these ultra-large models. The common feature among them is their rich prior knowledge, understanding, and OOD capabilities, which form the foundation of training-free missing multi-modality completion.

#### **4.3.** Conditional Generation

Conditional generation creates new data samples based on specific inputs, enabling models to generate tailored outputs. Early foundational work in this area includes Generative Adversarial Networks (GANs) [11] and Variational Autoencoders (VAEs) [19], with notable advancements such as conditional GANs [32] and Conditional VAEs (CVAEs) [42]. Recent progress has introduced diffusion models [23, 41, 43], which have demonstrated the ability to generate more realistic content, as exemplified by systems like DALL-E [39] and Stable Diffusion [36]. ControlNet [56] further improves user control. The key distinction between these approaches and our method lies in their respective objectives. Unlike these approaches, which focus on diverse



Figure 3. **Visualization analysis.** We present the results of multi-modality completion across different datasets. These results include visualizations related to the number of objects and specific people cases in general domain, as well as lesion completion in the medical domain. The green box represents the part that is close to the ground-truth, and the red box represents the wrong part.

content generation, our method emphasizes understanding modality interactions to **accurately** complete **missing content**.

### 5. Discussion and Limitations

Why we need the training-free MMC? Previous approaches for handling missing modalities have demonstrated limited generalizability across different domains and constrained data flexibility. For example, recent promptbased methods such as MPMM [21] and MPLMM [13] have shown significantly lower performance in general domains compared to their effectiveness in the medical domain, as reflected in our experimental results. Additionally, these methods depend on pre-trained models that are highly tailored to the target domain. In data-hungry scenarios, such as the medical field where data often contain missing elements, these constraints greatly impact the reliability of model decisions. In recent years, LMMs have exhibited remarkable zero-shot capabilities across various domains. Consequently, leveraging these models in a training-free manner to address the MMC problem presents a promising and cost-effective research direction. Our experimental results corroborate this perspective, and we hope these findings will encourage the community to further explore the potential of LMMs for addressing MMC challenges.

Why 'knowledge bridger' can help MMC? We believe that the key to addressing the MMC problem using LMM lies in accurately generating missing data and effectively ranking the generation candidates. Our experimental results show that directly employing LMM to generate the missing modality does not ensure the accuracy required for MMC. Our proposed approach first involves mining internal knowledge from the available modalities, enabling LMM to understand intra-modal interactions. This structured knowledge is then used to address challenges in both generation and ranking. However, using LMM inherently leads to hallucinations. In future work, we plan to employ more robust knowledge extraction methods, such as retrieval-augmented generation [14], to mitigate hallucinations<sup>6</sup> and further enhance the accuracy of missing data generation.

Limitations. Our method focuses exclusively on image and text modalities, leaving its performance on other modalities, such as speech and depth, yet to be explored. The approach emphasizes the automatic extraction of inter-modal knowledge and the completion of missing modalities through domain knowledge. Thus, in the future, adaptation to other modalities is possible by defining a more comprehensive modality knowledge. Some promising works [10, 29] show that one modality, such as image or text, can be connected to any other modality. Additionally, we observe that while our method enhances classification performance under a high missing rate (e.g., 0.7), it paradoxically results in a decrease in the similarity scores of the completed modalities. Therefore, there remains substantial potential for further exploration to develop more robust generation and ranking strategies in the future.

#### 6. Conclusion

In this paper, we propose a novel MMC challenge and introduce a training-free MMC pipeline called 'Knowledge Bridger' based on LMM. Our findings demonstrate that LMM can autonomously mine knowledge across modalities and leverage this structured knowledge to enhance the accuracy of missing modality generation and ranking. Furthermore, experimental results indicate that as the model's parameter scale increases, our method shows improved knowledge extraction capabilities.

<sup>&</sup>lt;sup>6</sup>Potential hallucinations of our method are presented in appendix.

Acknowledgment. This work is supported by the Guangdong Natural Science Funds for Distinguished Young Scholars (Grant 2023B1515020097), the AI Singapore Programme under the National Research Foundation Singapore (Grant AISG3-GV-2023-011), the Beijing Natural Science Foundation (No.4234086), the Natural Science Foundation of China (No. 62192782), and the Lee Kong Chian Fellowships.

### References

- John Arevalo, Thamar Solorio, Manuel Montes-y Gómez, and Fabio A González. Gated multimodal units for information fusion. In *ICLR*, 2017. 4
- [2] Simone Bianco, Luigi Celona, Marco Donzella, and Paolo Napoletano. Improving image captioning descriptiveness by ranking and llm-based fusion. arXiv preprint arXiv:2306.11593, 2023. 4
- [3] Lei Cai, Zhengyang Wang, Hongyang Gao, Dinggang Shen, and Shuiwang Ji. Deep adversarial learning for multimodality missing data completion. In *SIGKDD*, pages 1158– 1166, 2018. 1
- [4] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, pages 104–120. Springer, 2020. 7
- [5] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *CVPR*, pages 24185–24198, 2024. 7
- [6] Noah Cohen Kalafut, Xiang Huang, and Daifeng Wang. Joint variational autoencoders for multimodal imputation and embedding. *Nature Machine Intelligence*, 5(6):631–642, 2023. 1
- [7] Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310, 2016. 4, 6
- [8] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. In *ICML*, 2023. 7
- [9] Georgios Evangelopoulos, Athanasia Zlatintsi, Alexandros Potamianos, Petros Maragos, Konstantinos Rapantzikos, Georgios Skoumas, and Yannis Avrithis. Multimodal saliency and fusion for movie summarization based on aural, visual, and textual attention. *IEEE TMM*, 15(7):1553–1568, 2013. 7
- [10] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In CVPR, pages 15180–15190, 2023. 8

- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 7
- [12] Shuhao Gu, Jialing Zhang, Siyuan Zhou, Kevin Yu, Zhaohu Xing, Liangdong Wang, Zhou Cao, Jintao Jia, Zhuoyi Zhang, Yixuan Wang, et al. Infinity-mm: Scaling multimodal performance with large-scale and high-quality instruction data. arXiv preprint arXiv:2410.18558, 2024. 5
- [13] Zirun Guo, Tao Jin, and Zhou Zhao. Multimodal prompt learning with missing modalities for sentiment analysis and emotion recognition. In ACL, pages 1726–1736, 2024. 1, 4, 6, 7, 8
- [14] Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. Lightrag: Simple and fast retrieval-augmented generation. arXiv preprint arXiv:2410.05779, 2024. 8
- [15] Jing Han, Zixing Zhang, Zhao Ren, and Björn Schuller. Implicit fusion by joint audiovisual training for emotion recognition in mono modality. In *ICASSP*, pages 5861–5865. IEEE, 2019. 7
- [16] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. arXiv preprint arXiv:2410.21276, 2024. 1, 2, 5, 7
- [17] Xin Jiang, Hao Tang, Junyao Gao, Xiaoyu Du, Shengfeng He, and Zechao Li. Delving into multimodal prompting for fine-grained visual classification. In AAAI, pages 2570– 2578, 2024. 1
- [18] Guanzhou Ke, Guoqing Chao, Xiaoli Wang, Chenyang Xu, Yongqi Zhu, and Yang Yu. A clustering-guided contrastive fusion for multi-view representation learning. *IEEE T-CSVT*, 34(4):2056–2069, 2023. 7
- [19] Diederik P Kingma. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013. 7
- [20] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *NIPS*, 35:22199–22213, 2022. 1, 3
- [21] Yi-Lun Lee, Yi-Hsuan Tsai, Wei-Chen Chiu, and Chen-Yu Lee. Multimodal prompting with missing modalities for visual recognition. In *CVPR*, pages 14943–14952, 2023. 1, 4, 6, 7, 8
- [22] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, pages 19730–19742. PMLR, 2023. 2, 4
- [23] Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-Im improves controllable text generation. In NIPS, pages 4328–4343, 2022. 7
- [24] Zheng Lian, Lan Chen, Licai Sun, Bin Liu, and Jianhua Tao. Gcnet: Graph completion network for incomplete multimodal learning in conversation. *IEEE T-PAMI*, 45(7):8419– 8432, 2023. 7
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014. 4, 6
- [26] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *NIPS*, 36, 2024. 1, 7

- [27] Shanglin Liu, Jianming Lv, Jingdan Kang, Huaidong Zhang, Zequan Liang, and Shengfeng He. Modfinity: Unsupervised domain adaptation with multimodal information flow intertwining. In CVPR, 2025. 1
- [28] Jianming Lv, Kaijie Liu, and Shengfeng He. Differentiated learning for multi-modal domain adaptation. In ACM MM, pages 1322–1330, 2021. 1
- [29] Yuanhuiyi Lyu, Xu Zheng, Jiazhou Zhou, and Lin Wang. Unibind: Llm-augmented unified and balanced representation space to bind them all. In *CVPR*, pages 26752–26762, 2024. 8
- [30] Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. Smil: Multimodal learning with severely missing modality. In AAAI, pages 2302–2310, 2021. 1, 7
- [31] Mengmeng Ma, Jian Ren, Long Zhao, Davide Testuggine, and Xi Peng. Are multimodal transformers robust to missing modality? In CVPR, pages 18177–18186, 2022. 7
- [32] Mehdi Mirza. Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784, 2014. 7
- [33] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *ICML*, pages 689–696, 2011. 1
- [34] Srinivas Parthasarathy and Shiva Sundaram. Training strategies to handle missing modalities for audio-visual expression recognition. In *ICMI*, pages 400–404, 2020. 7
- [35] Hai Pham, Paul Pu Liang, Thomas Manzini, Louis-Philippe Morency, and Barnabás Póczos. Found in translation: Learning robust joint representations by cyclic translations between modalities. In AAAI, pages 6892–6899, 2019. 7
- [36] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdx1: Improving latent diffusion models for high-resolution image synthesis. In *ICLR*, 2024. 5, 7
- [37] S Prabu, M Lakshmanan, and V Noor Mohammed. A multimodal authentication for biometric recognition system using intelligent hybrid fusion techniques. *Journal of medical systems*, 43(8):249, 2019. 7
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 2, 4, 5
- [39] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, pages 8821– 8831. Pmlr, 2021. 7
- [40] Sucheng Ren, Yong Du, Jianming Lv, Guoqiang Han, and Shengfeng He. Learning from the master: Distilling crossmodal advanced knowledge for lip reading. In *CVPR*, pages 13325–13333, 2021. 1
- [41] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In CVPR, pages 22500–22510, 2023. 7
- [42] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *NIPS*, 28, 2015. 7

- [43] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2020. 7
- [44] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024.1, 7
- [45] Gemini Team. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv preprint arXiv:2403.05530, 2024. 7
- [46] Hu Wang, Yuanhong Chen, Congbo Ma, Jodie Avery, Louise Hull, and Gustavo Carneiro. Multi-modal learning with missing modality via shared-specific feature modelling. In *CVPR*, pages 15878–15887, 2023. 7
- [47] Junke Wang, Lingchen Meng, Zejia Weng, Bo He, Zuxuan Wu, and Yu-Gang Jiang. To see is to believe: Prompting gpt-4v for better visual instruction tuning. arXiv preprint arXiv:2311.07574, 2023. 5
- [48] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. arXiv preprint arXiv:2409.12191, 2024. 1, 5, 7
- [49] Yuanzhi Wang, Zhen Cui, and Yong Li. Distributionconsistent modal recovering for incomplete multimodal learning. In *ICCV*, pages 22025–22034, 2023. 1, 4, 6
- [50] Zilong Wang, Zhaohong Wan, and Xiaojun Wan. Transmodality: An end2end fusion method with transformer for multimodal sentiment analysis. In WWW, pages 2514–2520, 2020. 7
- [51] Tobias Weber, Michael Ingrisch, Bernd Bischl, and David Rügamer. Cascaded latent diffusion models for highresolution chest x-ray synthesis. In *Pacific-Asia Conference* on Knowledge Discovery and Data Mining, pages 180–191. Springer, 2023. 5
- [52] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. arXiv preprint arXiv:2206.07682, 2022. 2
- [53] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. NPIS, 35:24824–24837, 2022. 2
- [54] Jiandian Zeng, Tianyi Liu, and Jiantao Zhou. Tag-assisted multimodal sentiment analysis under uncertain missing modalities. In SIGIR, pages 1545–1554, 2022. 1, 7
- [55] Chaohe Zhang, Xu Chu, Liantao Ma, Yinghao Zhu, Yasha Wang, Jiangtao Wang, and Junfeng Zhao. M3care: Learning with missing modalities in multimodal healthcare data. In *SIGKDD*, pages 2418–2428, 2022. 1
- [56] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *CVPR*, pages 3836–3847, 2023. 7
- [57] Jinming Zhao, Ruichen Li, and Qin Jin. Missing modality imagination network for emotion recognition with uncertain missing modalities. In ACL, pages 2608–2618, 2021. 1, 4, 6, 7